arXiv:1403.7560v1 [physics.comp-ph] 28 Mar 2014

CUDA programs for GPU computing of Swendsen-Wang multi-cluster spin flip algorithm: 2D and 3D Ising, Potts, and XY models

Yukihiro Komura[a,b], Yutaka Okabe[b]

[a] *Nuclear Science and Engineering Directorate, Japan Atomic Energy Agency, 2-4 Shirakata-shirane, Tokai-mura, Naka-gun, Ibaraki 319-1195, Japan*
[b] *Department of Physics, Tokyo Metropolitan University, Hachioji, Tokyo 192-0397, Japan*

Abstract

We present sample CUDA programs for the GPU computing of the Swendsen-Wang multi-cluster spin flip algorithm. We deal with the classical spin models; the Ising model, the q-state Potts model, and the classical XY model. As for the lattice, both the 2D (square) lattice and the 3D (simple cubic) lattice are treated. We already reported the idea of the GPU implementation for 2D models [Comput. Phys. Commun. 183 (2012) 1155-1161]. We here explain the details of sample programs, and discuss the performance of the present GPU implementation for the 3D Ising and XY models. We also show the calculated results of the moment ratio for these models, and discuss phase transitions.

Keywords: Monte Carlo simulation, cluster algorithm, Ising model, XY model, parallel computing, GPU

Program summary

Program title: CUDA programs for GPU computing of Swendsen-Wang multi-cluster spin flip algorithm

Catalogue identifier: AAAA

Program summary URL: http://cpc.cs.qub.ac.uk/summaries/AAAA.html

Program obtainable from: CPC Program Library, Queen's University, Belfast,

Email addresses: komura.yukihiro@jaea.go.jp (Yukihiro Komura), okabe@phys.se.tmu.ac.jp (Yutaka Okabe)

N. Ireland

Licensing provisions: Standard CPC licence, http://cpc.cs.qub.ac.uk/licence/licence.html

No. of lines in distributed program, including test data, etc.: ****

No. of bytes in distributed program, including test data, etc.: ****

Distribution format: tgz

Programming language: C, CUDA

Computer: System with an NVIDIA CUDA enabled GPU

External routines: NVIDIA CUDA Toolkit 3.0 or newer

Nature of problem: Monte Carlo simulation of classical spin systems. Ising, q-state Potts model, and the classical XY model are treated for both two-dimensional and three-dimensional lattices.

Solution method: GPU-based Swendsen-Wang multi-cluster spin flip Monte Carlo method. The CUDA implementation for the cluster-labeling is based on the work by Hawick *et al.* [K.A. Hawick, A. Leist, and D. P. Playne, Parallel Computing 36 (2010) 655-678], and that by Kalentev *et al.* [O. Kalentev, A. Rai, S. Kemnitzb, and R. Schneider, J. Parallel Distrib. Comput. 71 (2011) 615-620].

Restrictions: The system size is limited depending on the memory of a GPU.

Running time: For the parameters used in the sample programs, it takes about a minute for each program. Of course, it depends on the system size, the number of Monte Carlo steps, etc.

1. Introduction

Monte Carlo simulation is a standard method to study statistical physics of many body problems [1]. The advance in computer hardware together with the development of software has enhanced the importance of Monte Carlo simulation all the time. The general purpose computing using the graphics processing unit (GPU) is a hot topic in computer science; using the common unified device architecture (CUDA) released by NVIDIA, it is now easy to implement algorithms on GPU using standard C or C++ language with CUDA specific extension. Preis *et al.* [2, 3] proposed the efficient use of GPU for the Metropolis-type [4] single spin

flip Monte Carlo simulation of classical spin systems. Preis *et al.* [2] reported the acceleration of the computation 60 times for the two-dimensional (2D) Ising model, and 35 times for the three-dimensional (3D) Ising model compared to the CPU computation using a current CPU core. The parallel computing using a GPU is realized by interpenetrating sublattice decomposition, which is the same as vector computers in 1980s. Block *et al.* [5] further discussed the use of multiple GPUs for the Monte Carlo simulation as well as the multispin coding technique.

The single spin flip Metropolis simulation often suffers from the problem of slow dynamics; that is, it takes long time for equilibration, for example, at temperatures near the critical temperature of the phase transition. To conquer the problem of slow dynamics, the cluster spin flip algorithms of Monte Carlo simulation have been proposed. The multi-cluster spin flip algorithm due to Swendsen and Wang [6] and the single-cluster spin flip algorithm due to Wolff [7] are typical examples. The parallelization of cluster spin flip algorithm is not straightforward because the cluster labeling part of the cluster spin flip algorithm basically requires a sequential calculation, which is in contrast to the local calculation for the single spin flip algorithm. The present authors [8] proposed the GPU computing for the Swendsen-Wang (SW) multi-cluster spin flip algorithm, where the idea of Hawick *et al.* [9] was used in the cluster labeling part. An improved version of the cluster labeling algorithm was proposed by Kalentev *et al.* [10]. The present authors [8] reported the acceleration of the GPU computing with 12.4 times for the SW algorithm of the 2D Ising model, in comparison with the CPU computing using a current CPU core. The present authors [11] extended the GPU computation for the SW multi-cluster spin flip algorithm to the multiple GPU computation, where two steps of parallelization are used, that is, the parallelization within each GPU and the inter-GPU parallelization. The multiple GPU computation of the multi-cluster spin flip algorithm was applied to the large-scale Monte Carlo simulation of the 2D classical XY model up to the system size 65536×65536 [12]; the Kosterlitz-Thouless [13] transition temperature and the exponent to specify the multiplicative loga-

rithmic correction were precisely determined. We note that the present authors also proposed the GPU computing for the Wolff single-cluster spin flip algorithm [14], although the efficiency of the acceleration was not so good as compared to the case of the multi-cluster spin flip algorithm.

Here, we publicize sample CUDA programs for the GPU computing of the multi-cluster spin flip algorithm, and explain the details of the sample programs. We deal with the following classical spin models;

- 2D Ising model,
- 2D Potts model,
- 2D classical XY model,
- 3D Ising model,
- 3D Potts model,
- 3D classical XY model.

As for the lattice, both the 2D (square) lattice and the 3D (simple cubic) lattice are treated.

The Hamiltonian of the Ising model is given by

$$\text{(Ising)} \quad \mathcal{H} = -J \sum_{\langle i,j \rangle} s_i s_j, \quad s_i = \pm 1. \tag{1}$$

Here, J is the coupling and s_i is the Ising spin on the lattice site i. The summation is taken over the nearest neighbor pairs $\langle i, j \rangle$. Periodic boundary conditions are employed. The q-state Potts model is an extension of the Ising model to the Potts spin with q states, whose Hamiltonian is

$$\text{(Potts)} \quad \mathcal{H} = -J \sum_{\langle i,j \rangle} (\delta_{s_i,s_j} - 1), \quad s_i = 1, 2, \cdots, q, \tag{2}$$

where $\delta_{a,b}$ is the Kronecker's delta. The 2-state Potts model is essentially the same as the Ising model except for the energy unit, but we separately give the program of the Ising model because it is easy for readers to start with the Ising

model as a basic model. Finally, the Hamiltonian of the classical XY model, a continuous spin model, is given by

$$\text{(classical XY)} \quad \mathcal{H} = -J \sum_{\langle i,j \rangle} \boldsymbol{s}_i \cdot \boldsymbol{s}_j = -J \sum_{<i,j>} \cos(\theta_i - \theta_j), \tag{3}$$

where $\boldsymbol{s}_i$ is a planar unit vector with two real components, $(\cos\theta_i, \sin\theta_i)$. For actual implementation, the value of θ_i is discretized as $2\pi p_i/q$ with $p_i = 1, 2, \cdots, q$. This discretized model is referred to as the q-state clock model. When q tends to infinity, the clock model becomes the classical XY model. To make a cluster flip of vector spins, we use the idea of embedded cluster introduced by Wolff [7].

The rest of the paper is organized as follows. In section 2, we briefly describe the SW multi-cluster spin flip algorithm and its GPU implementation for the Ising model as an example. In section 3, we report the performance of the present GPU implementation for the 3D Ising and XY models. We show the calculated results of the moment ratio for these models, and discuss the phase transitions. We note that we already reported the performance of the 2D models in Ref. [8]. We give the summary and discussion in section 4.

2. GPU computation of Swendsen-Wang cluster algorithm

In the SW multi-cluster spin flip algorithm [6], we flip clusters of spins at once. The SW algorithm for the Ising model, Eq. (1), consists of three main steps:

(1) *Step of active bond generation*: Construct a bond lattice of active or non-active bonds with probability $p = 1 - e^{2J/T}$, where T is the temperature.

(2) *Step of cluster labeling*: The active bonds partition the spins into clusters which are identified and labeled using a cluster-labeling algorithm.

(3) *Step of spin flip*: All spins in each cluster are set randomly to +1 or -1.

We should consider how to make parallel computation when we use GPUs. Since the calculations of the step of active bond generation and the step of spin flip are done independently on each site, these steps are well suited for parallel

computation on GPU. In the CPU computation of the step of cluster labeling, the Hoshen-Kopelman algorithm [15], which was first introduced in context of cluster percolation, is often used as an efficient algorithm. The Hoshen-Kopelman algorithm is a special version of the class of union-and-find algorithms [16]. Since the assignment of label of cluster is done on each site piece by piece sequentially in the Hoshen-Kopelman algorithm, it is not straightforward to extend to parallel computing.

Recently, Hawick *et al.* [9] investigated the cluster-labeling algorithm efficient for GPU calculation. They proposed the labeling method of "Label Equivalence". The procedure of their algorithm was schematically explained in Fig. 1 of Ref. [8]. In the CUDA program, the GPU "kernel" is invoked. The algorithm of Hawick *et al.* consists of three kernel functions, that is, scanning function, analysis function and labeling function, and two variables for labeling; one is a variable for saving the label, `label`, and the other is a temporal variable for updated label, `R`. The scanning function compares the label of each site with that of the nearest-neighbor sites when the bond between each site and the nearest-neighbor site is active. If the label of the nearest-neighbor site is smaller than the label of that site, the temporal variable with the label number, `R[label[index]]`, is updated to the smallest one. For the update of the temporal variable on the scanning function, the atomic operation `atomicMin()` is used. Atomic operations provided by CUDA are performed without interference from any other threads. The analysis function resolves the equivalence chain of `R`; the temporal variable `R[index]` is updated from the starting site to the new site. Each thread checks the temporal variable and the label on each site. When the label number, `label`, is equal to the thread number, `index`, each thread tracks back the temporal variable until the temporal variable, `R`, remains unchanged. Since each thread executes this operation concurrently, the final value is reached quickly. The labeling function updates the label for saving by `label[index]` $\leftarrow$ `R[label[index]]`. In the cluster-labeling algorithm of Hawick *et al.*, the loop including three functions is iterated up to the point when the information of the labeling needs no more process of scanning function.

More recently, Kalentev *et al.* [10] reported the refinement of the algorithm. The procedure of their algorithm is explained in Fig. 2 of Ref. [8]. They used only one variable for labeling instead of two because there is no need for a temporal reference; the implementation was improved in terms of memory consumption. It means that the number of kernel functions are reduced from three to two because the process of the labeling function is no more needed. They changed the execution condition on the analysis function, and eliminated the atomic operation. With the refinements due to Kalentev *et al.*, the improvement of computational speed and the reduction of the memory usage were realized.

Now we explain our program `AAAA` in detail. The system size is represented by `nx*nx` for 2D models and `nx*nx*nx` for 3D models, where `nx` should be a multiple of 32. As for 2D models, the maximum size of `nx` is 4096 or so, depending on memory size of GPU. If we use GTX580, the maximum linear size is 4096, whereas that is 8192 if we use GTX680, for example. As for 3D models, the maximum linear size is 256 for GTX580, whereas that is 512 for GTX680. For the cluster-labeling algorithm, we can choose either the algorithm of Hawick *et al.* (`algorithm = 0`), or the algorithm by Kalentev *et al.* (`algorithm = 1`). The number of Monte Carlo steps (MCSs) per spin for discard is denoted by `mcs1`, whereas the number of MCSs per spin for measurement is denoted by `mcs2`.

The kernel function

$$device_function_spinset;$$

is a function for setting initial spin configuration. The kernel function for the step of active bond generation is

$$device_function_init_H;$$

and

$$device_function_init_K;$$

for the algorithm of Hawick *et al.* and for the algorithm of Kalentev *et al.*, respectively. Three kernel functions,

$$device_function_scanning_H;$$
$$device_function_analysis_H;$$
$$device_function_labeling_H;$$

are used in the step of cluster labeling for the algorithm of Hawick *et al.*, whereas two kernel functions

$$device_function_scanning_K;$$
$$device_function_analysis_K;$$

are used in the step of cluster labeling for the algorithm of Kalentev *et al.* The kernel functions

$$device_function_spin_select;$$
$$device_function_spin_flip;$$

are used in the step of spin flip. The kernel function

$$device_function_sum;$$

is a function for measuring physical quantities.

For this sample program, we measure the magnetization

$$M = \sum_i s_i, \tag{4}$$

and the total energy, Eq. (1). In the calculation of total sum, reduction of data is important to avoid data confliction. For the outputs, we print the 2nd moment of magnetization, the 4th moment of magnetization, the total energy per spin, and the specific heat per spin. We also print the computational time for spin-flip and the computational time for spin-flip together with measurement of magnetization and energy. The computation time is represented in units of milliseconds per a single MCS. It is easy to extend our sample program to

calculate other quantities, such as the correlation function, the helicity modulus, etc.

We make comments on the technical part. The proper use of shared memories is important to save computational time. To improve the computational speed and save memory, we store the information on spin, bond and label in one word. Here, we use a linear congruential random-number generator which was proposed by Preis *et al.* [2], but any other random-number generators can be used.

A brief remark is made on the q-state Potts model. The program of the Potts model is essentially the same as that of the Ising model except that the number of states is q instead of 2. The square of the magnetization of the q-state Potts model is calculated as

$$M^2 = \frac{q \sum_{k=1}^{q} n[k]^2 - N^2}{q-1}, \tag{5}$$

where $n[k]$ is the number of spins with the state k, and N is the total number of spins.

Next, we discuss the GPU-based calculation of SW multi-cluster algorithm for the XY model, actually the q-state clock model. To make a cluster flip, we use the idea of embedded cluster introduced by Wolff [7]. We project vector spins to form Ising spin clusters. The essential part of the GPU implementation is the same as the case of the Ising model. The variable `ispr`, which is randomly chosen, specifies the direction for mirror projection, and it is constant for all the sites, that is, all the threads. We note that the proper use of shared memories is effective especially for the calculation of the inner product of vectors. The number of states q should be less than or equal to 512 because we use shared memory for the table of rule for energy, `rule[ ]`, and the table for the magnetization, `fmy[ ]`.

3. Results

We have tested the performance of our code on NVIDIA GeForce GTX580. Since the performance for 2D models was discussed in Ref. [8], we here show

the data for 3D models.

		L=32	L=64	L=128	L=256
GTX580	update only	7.59 nsec	3.52 nsec	3.01 nsec	3.10 nsec
Kalentev *et al.*	+ measurement	8.87 nsec	3.90 nsec	3.31 nsec	3.33 nsec

Table 1: Average computational time per a spin flip for the 3D Ising model at $T = 4.5115$. The time for only a spin update and that including the measurement of energy and magnetization are given.

We first deal with the 3D Ising model. For the cluster-labeling algorithm, we show the data using the algorithm due to Kalentev *et al.* [10], because it is always faster than the calculation using that due to Hawick *et al.* [9]. We calculated the 3D Ising model near the critical temperature, $T = 4.5115$, [19, 20]. Throughout the present paper, we measure the temperature in units of J. The average computational times per a spin update near the critical temperature for the Ising model are tabulated in Table 1. There, the time for only a spin update and that including the measurement of energy and magnetization are given. We show the measured time per a spin flip in units of nano sec. The linear system sizes L are L=32, 64, 128 and 256. We can see from Table 1 that the computational time of our GPU implementation of the SW algorithm is almost constant for $L \geq 128$; for smaller lattices the parallelization is not so effective. The performance for the 2D Ising model with the algorithm of Kalentev *et al.* was 2.51 nano sec per a spin flip with $L = 4096$ on GTX580 [8]. It means that the computational time for a spin flip of the 3D model is compatible with that of the 2D model although the nearest neighbor sites are 6 instead of 4 for the simple cubic lattice.

We next discuss the temperature dependence of physical quantities of the 3D Ising model. It is known that this model undergoes the 2nd-order phase transition. The linear system sizes are $L = 32$, 64, 128, 256, and 512. We discarded the first 10,000 MCSs, and the next 200,000 MCSs were used for measurement for $L \leq 128$. The MCSs for measurement were 100,000 and 50,000 for $L = 256$ and for $L = 512$, respectively. We used GeForce GTX680 for

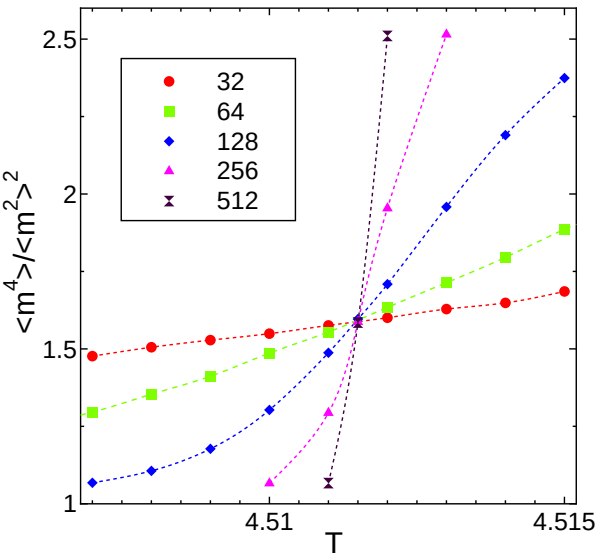

Figure 1: Moment ratio of 3D Ising model; L=32, 64, 128, 256, and 512.

$L = 512$. We plot the moment ratio

$$U(T) = \frac{\langle M(T)^4 \rangle}{\langle M(T)^2 \rangle^2} = \frac{\langle m(T)^4 \rangle}{\langle m(T)^2 \rangle^2} \tag{6}$$

with $m = M/N$, which is essentially the Binder ratio [18] except for the normalization, of the 3D Ising model in Figure 1. We made four independent runs for each size; the average was taken over four runs, and the statistical errors were estimated, which are smaller than the size of the marks in Figure 1. The crossing of the data with different sizes reproduces the known result of the critical temperature.

Let us consider the finite-size scaling (FSS) [17] of the moment ratio $U(T)$. We plot $U(T) = \langle m^4 \rangle / \langle m^2 \rangle^2$ as a function of $tL^{1/\nu}$; all the data with different sizes are collapsed on a single curve. Here, $t = (T - T_c)/T_c$, and ν is the critical exponent for the correlation length. With the proper choice of T_c and ν, we can estimate T_c and ν. The FSS plot for the 3D Ising model is given in Figure 2. The FSS is extremely good; a wrong value of the estimated T_c yields a large deviation of FSS for large sizes. That is, we can get precise estimates with the data of

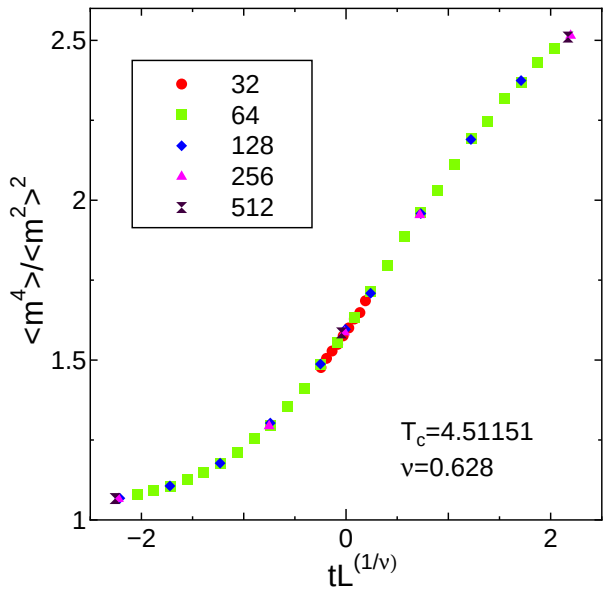

Figure 2: FSS plot of moment ratio of 3D Ising model, where $t = (T - T_c)/T_c$ with T_c=4.51151 and ν=0.628.

large sizes. Here, the best choices of T_c and ν are 4.51151(1) and 0.628(4), respectively. The previous estimates of T_c are 4.511424(53) (=1/0.2216595(26)) (Ref. [19]) and 4.511524(20) (=1/0.2216546(10)) (Ref. [20]). We conclude that our estimate using large sizes up to 512 is consistent with that of [20] within the statistical errors, whereas it is outside of the error bars from that of [19]. The detailed analysis will be left to a future work, where much larger sizes will be treated with multiple GPUs.

		L=32	L=64	L=128	L=256
GTX580	update only	8.78 nsec	4.62 nsec	4.15 nsec	3.96 nsec
Kalentev *et al.*	+ measurement	10.33 nsec	5.27 nsec	4.57 nsec	4.49 nsec

Table 2: Average computational time per a spin flip for the 3D XY model at $T = 2.20175$. The time for only a spin update and that including the measurement of energy and magnetization are given.

The average computational times per a spin update near the critical temperature, $T = 2.20175$, for the classical XY model are tabulated in Table 2.

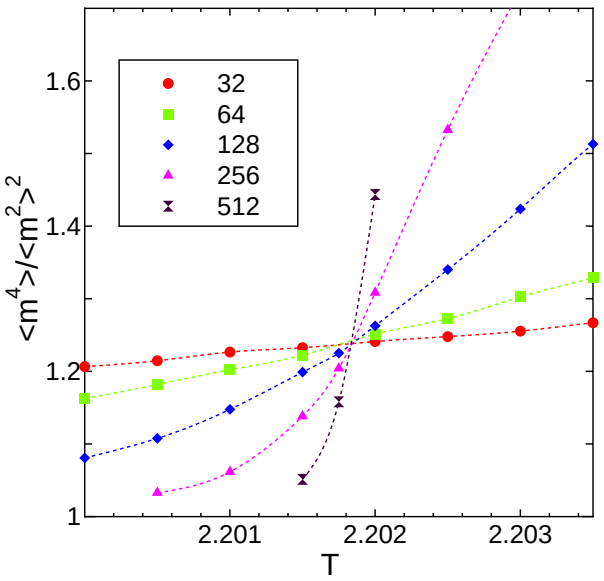

Figure 3: Moment ratio of 3D XY model; L=32, 64, 128, 256, and 512.

The previous estimates of the critical temperature were given in [21, 22, 23]. Actually, we dealt with the 512-state clock model. The discritization level is small enough except for very low temperatures. Again, the time for only a spin update and that including the measurement of energy and magnetization are given. We show the measured time per a spin flip in units of nano sec. The linear system sizes L are L=32, 64, 128 and 256. We can see from Table 2 that the computational time of our GPU implementation of the SW algorithm is almost constant for $L \geq 128$. The computational time for the XY model is longer than the case of the Ising model because it needs more calculation for the projection to Ising spins. But such an extra computational time is only a little, which shows that our GPU implementation of the SW multi-cluster algorithm is also effective for the vector spin model, the clock model.

We plot the moment ratio of the 3D XY model in Figure 3. The linear system sizes are L = 32, 64, 128, 256, and 512. The number of the MCSs for discard and measurement are the same as the case of the Ising model. We made

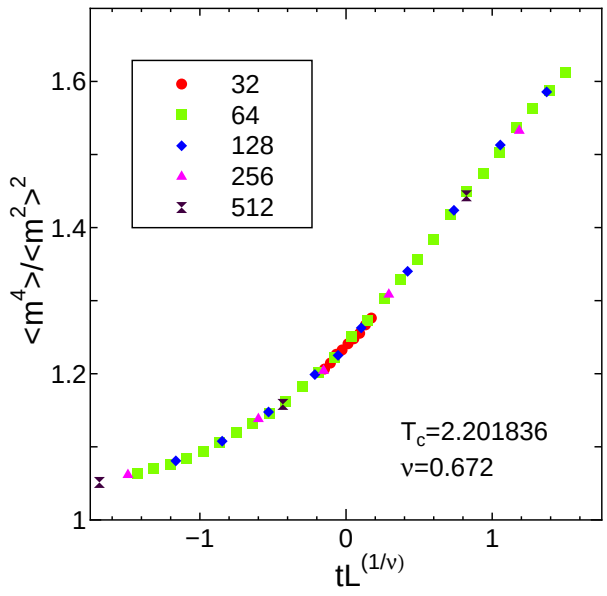

Figure 4: FSS plot of moment ratio of 3D XY model, where $t = (T - T_c)/T_c$ with T_c=2.201836 and ν=0.672.

four independent runs for each size; the average was taken over four runs. The crossing of the data with different sizes again reproduces the known results of the critical temperatures.

We show the FSS plot of $U(T) = \langle m^4 \rangle / \langle m^2 \rangle^2$ as a function of $tL^{1/\nu}$ in Fig. 4. Here, the best choices of T_c and ν are 2.201836(6) and 0.672(4), respectively. The authors of recent publication [23] gave two estimates of T_c as 2.201852(1) and 2.2018312(6). Our estimate is between the two. The detailed analysis of the 3D classical XY model will also be left to a future work, where much larger sizes will be treated with multiple GPUs.

4. Summary and discussion

We have explained our program for the GPU-based SW multi-cluster spin flip algorithm for the classical spin models. Treated models are the Ising model, the q-state Potts model, and the classical XY model for 2D and 3D lattices. For the assignment of clusters, we used the algorithm by Hawick *et al.* [9] and that

by Kalentev *et al.* [10]. We have shown the performance of the 3D Ising model and the 3D classical XY model. The performance of GPU calculation is good, which is the same as the 2D case [8]. Calculating the moment ratios of the 3D Ising and classical XY models, the critical temperature T_c and the correlation-length exponent ν are discussed by using the FSS analysis. The system sizes we treated are up to $L = 512$ for 3D systems, but much larger sizes can be treated if we use multiple GPUs [11]. In near future, we make mutiple-GPU calculations for getting precise estimates for 3D models.

In this paper, we have shown sample programs for the GPU-based SW multi-cluster spin flip algorithm of classical XY models. For the physical quantities we only calculated the magnetization and the energy. The extension to the calculation of correlation function, helicity modulus, etc. is straightforward. We can extend the GPU calculation to other cluster algorithms including the quantum Monte Carlo simulations. We hope more researchers show interest in the GPU calculations of the cluster algorithms.

Acknowledgment

This work was supported by a Grant-in-Aid for Scientific Research from the Japan Society for the Promotion of Science.

References

[1] D.P. Landau, K. Binder, *A Guide to Monte Carlo Simulations in Statistical Physics*, 3rd edition (Cambridge University Press, Cambridge, 2009)

[2] T. Preis, P. Virnau, W. Paul, J.J. Schneider, GPU accelerated Monte Carlo simulation of the 2D and 3D Ising model, J. Comp. Phys. 228 (2009) 4468-4477.

[3] B. Block, T. Preis, Computer simulations of the Ising Model on Graphics Processing Units, Eur. Phys. J. Special Topics 210 (2012) 133-145.

[4] N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, E. Teller, Equation of State Calculations by Fast Computing Machines, J. Chem. Phys. 21 (1953) 1087-1092.

[5] B. Block, P. Virnau, T. Preis, Multi-GPU accelerated multi-spin Monte Carlo simulations of the 2D Ising model, Comp. Phys. Comm. 181 (2010) 1549-1556.

[6] R.H. Swendsen, J.S. Wang, Nonuniversal critical dynamics in Monte Carlo simulations, Phys. Rev. Lett. 58 (1987) 86-88.

[7] U. Wolff, Collective Monte Carlo Updating for Spin Systems, Phys. Rev. Lett. 62 (1989) 361-364.

[8] Y. Komura, Y. Okabe, GPU-based Swendsen-Wang multi-cluster algorithm for the simulation of two-dimensional classical spin systems, Comput. Phys. Comm. 183 (2012) 1155-1161.

[9] K.A. Hawick, A. Leist, D. P. Playne, Parallel Graph Component Labelling with GPUs and CUDA, Parallel Computing 36 (2010) 655-678.

[10] O. Kalentev, A. Rai, S. Kemnitzb, R. Schneider, Connected component labeling on a 2D grid using CUDA, J. Parallel Distrib. Comput. 71 (2011) 615-620.

[11] Y. Komura, Y. Okabe, Multi-GPU-based Swendsen-Wang multi-cluster algorithm for the simulation of two-dimensional q-state Potts model Comput. Phys. Comm. 184 (2013) 40-44.

[12] Y. Komura, Y. Okabe, Large-scale Monte Carlo simulation of two-dimensional classical XY model using multiple GPUs, J. Phys. Soc. Jpn. 81 (2012) 113001.

[13] J. M. Kosterlitz, D. J. Thouless, Ordering, metastability and phase transitions in two-dimensional systems, J. Phys. C: Solid State Physics 6 (1973) 1181-1203.

[14] Y. Komura, Y. Okabe, GPU-based single-cluster algorithm for the simulation of the Ising model, J. Comput. Phys. 231 (2012) 1209-1215.

[15] J. Hoshen, R. Kopelman, Percolation and cluster distribution. I. Cluster multiple labeling technique and critical concentration algorithm, Phys. Rev. B 14 (1976) 3438-3445.

[16] T. H. Cormen, C. E. Leiserson, R. L. Rivest, C. Stein, Introduction to Algorithms, MIT Press, third edition, 2009.

[17] M. E. Fisher, in *Proc. 1970 E. Fermi Int. School of Physics*, edited by M. S. Green (Academic, New York, 1971) Vol. 51, pp. 1-99.

[18] K. Binder, Finite size scaling analysis of ising model block distribution functions, Z. Phys. B: Condens. Matter 43 (1981) 119-140.

[19] A.M. Ferrenberg, D.P. Landau, Critical behavior of the three-dimensional Ising model: A high-resolution Monte Carlo study, Phys. Rev. B 44 (1991) 5081-5091.

[20] H.W.J. Blöte, E. Luijten, J.R. Heringa, Ising universality in three dimensions: a Monte Carlo study, J. Phys. A: Math. Gen. 28 (1995) 6289-6314.

[21] A.P. Gottlob, M. Hasenbusch, S. Meyer, Critical behaviourof the 3D XY-model: A Monte Carlo study, Nucl. Phys. B (Proc. Suppl.) 30 (1993) 838-841.

[22] P. Butera, M. Comi, A.J. Guttmann, Critical exponents of the three-dimensional classical plane-rotator model on the sc lattice from a high-temperature-series analysis, Phys. Rev. B 48 (1993) 13987-13990.

[23] T.-Y. Lan, Y.-D. Hsieh, Y.-J. Kao, High-precision Monte Carlo study of the three-dimensional XY model on GPU, arXiv:1211.0780 (2012).